

ABSTRACT OF THE DISCLOSURE

An apparatus includes a computing node and a metadata server. The computing node may transmit a request to open a file to the metadata server, which may provide a corresponding file identifier to the computing node. The computing node may use the file identifier to directly access the storage storing the file. In one embodiment, the storage may be an object-based storage. The storage may receive the file identifier with an access command and may perform the mapping from file identifier to storage blocks internally.

Thus, accesses to the storage may be performed on an arbitrary boundary within the file. In other words, the entire block including the data needed by the computing node may not be transferred to the computing node. Instead, the storage may select the data requested from the file (e.g. via an offset and a number of bytes or some similar mechanism) and return that data to the computing node.